Audio/Video Inter-Device Power Control

Bruce Nordman, H.Y. Iris Cheung

Lawrence Berkeley National Laboratory

Abstract

There are many new digital communication technologies for procuring and displaying audio and video (A/V) content, and these are evolving rapidly in capability and integration into products. These enable new usage models for moving streams of A/V content among devices in a local network. A challenge this raises is ensuring that products do not remain fully on when they are not needed (often for hours or days between use) which can waste large amounts of energy. A near-term challenge is to create an architecture for how products should interact with respect to A/V content streams to maximize both user amenity and energy savings. Commonality across technologies and products is critical for this to succeed. In today's technologies, a content stream is either active or does not exist. This paper proposes a mechanism to address this problem – "sleeping streams". They have long-term persistence, going into a sleep state between uses, much as devices do when not in use. A stream that is asleep can be readily woken by any of the devices involved in it or by select external devices. This in turn should wake all devices needed for the stream to operate. When a device observes that it is involved with no active streams for a period of time, it can go to sleep. Knowledge of the stream and ability to control other devices, is distributed, not central. This paper addresses details of the concept, what it implies for existing and future A/V technologies, and the role of energy policy.

Introduction

Audio and video devices in U.S. homes are estimated to consume about 140 TWh/year [1]. The portion of this from devices on but not in use is not known, but even if it is only 10% of this, that is well over $1 billion/year in wasted energy. As the mechanisms described in this paper generally do not increase the manufacturing cost of products, attention to this topic is clearly merited.

Audio/Video systems today are of great variety in vintage, technologies, and size. While this paper primarily considers what power control of these devices should be like in future, with new devices, the current devices and context greatly influences that future in many ways.

A/V systems provide people with entertainment from traditional television broadcasts, subscription TV services, pre-recorded content, and, increasingly, from the Internet. These systems enable many different usage scenarios requiring different sets of devices to be on and active. With many devices, multiple users, and evolving technology, it is easy for devices to remain on long after the time they are needed, wasting energy. This paper proposes an architecture for managing A/V device power state with the key principle for each device to:

"wake up when it needs to, and go to sleep when it can".

This should enable power control to be more automatic and reliable, and provide greater usability to people as well as saving energy. The rest of this paper is organized as follows. First, key concepts for A/V systems are defined and explained. Then, the candidate architecture is proposed and detailed operational issues considered. Finally, there is consideration of the transition period with many legacy devices and a summary of next steps.

Audio/Video Systems

An A/V system is a collection of devices that provides video content (usually on a television), and/or audio content (usually on loudspeakers). Most A/V systems are found in residences, but they are present in all building types, including vehicles. Many A/V systems interact with devices in other locations via the Internet or in subscription television (cable/satellite) infrastructure. These remote devices are not part of the local A/V system. Information Technology (IT) devices like computers are increasingly connected to A/V systems, and to be compatible, including for power control, they need to implement the same A/V protocols and behaviors.

Typical residential A/V systems fall into two categories. In a simple system, there is no separate A/V receiver, so that all audio comes from the TV itself. Complex systems have a receiver which provides multi-channel sound and is used to select among the various sources (and can be a source itself). These systems commonly included Ethernet, Wi-Fi, HDMI, Audio, Composite video, and pay-TV connections. Many devices sold today already have some power management capabilities. For example, many subwoofers will go to sleep automatically if the audio input is silent for an extended period of time, and then wake when the audio reappears. Many A/V devices have an auto power-down feature that initiates after an extended time with no user input, but these typically lack a corresponding wake-up feature.

Concepts

The basic unit of an A/V system is a device, which has its own power state. Each device has a set of interfaces for communication. Connected interfaces on two devices form a link. The primary data transmitted across these links are streams of audio or video content. Devices can be the source (e.g. DVD player) or sink (e.g. TV or speaker) of a stream, or an intermediate device through which the stream passes that is not its source or sink. A component often ignored in considerations of A/V systems is the experience of the users who actually consume the content or otherwise interact with the devices. Data communicated other than the content streams includes control signals. Details of each of these concepts are covered below. Figure 1 shows how basic A/V devices map into the source, sink, and intermediate concepts. Figure 2 presents an abstract representation of the same thing.

Devices

Each device is individually powered, and have one or many communications interfaces, usually of a variety of technologies (see *Interfaces and Links* below), and has its own power state.

Power States

The overall power state of a device determines the types of functionality it is capable of, its responsiveness to the various interfaces it supports, its front panel user interface, and its power consumption. While the power levels of a device within a basic power state do vary somewhat, the basic power state is the primary determinant of overall energy use, and so the key consideration for saving energy in A/V devices.

Originally, A/V devices had only two basic power states: on and off. Decades later, two factors emerged to change this. First, the label “standby” was applied to any low-power state in which power consumption was non-zero, and as devices increasingly had off modes with non-zero power, standby became a synonym for off. Second, remote controls (usually with one-way infrared communication) were introduced which included the ability to turn a device on and off. When such devices were functionally off, they still required power to detect a power command from the remote control, creating a difference from the traditional off state. Such devices still had only two widely used power states, with disconnection from their power source occurring only rarely, and the “standby” state only exited with a power command and so still an off mode.

Some newer devices have three basic power states. An active mode, an off mode (which can only be exited with a power command), and an intermediate mode in which it is also responsive to communication from other devices. With Information Technology (IT) devices, these are usually called sleep states, but in A/V devices, no consistent terminology has been applied.

Interfaces and Links

An interface is hardware that enables a possible communications connection (link) to another device. Interfaces can be of many different wired technologies (e.g. composite video, line level audio, HDMI, or Ethernet) or, more recently, wireless.

Figure 1. Basic stream structure and typical components

Figure 2. Abstract stream structure (links to outside local network optional)

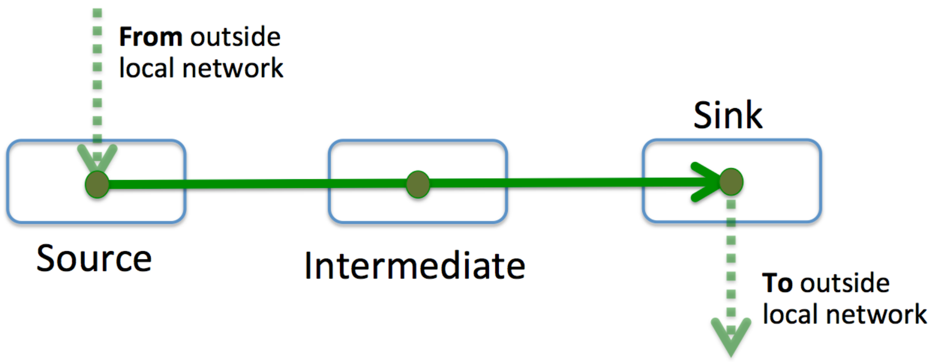

A data link is a point-to-point connection between two devices. There are many types: analog or digital, one-way or two-way, etc. A link is defined by the interfaces at each end, which must match. Originally, all A/V communication occurred in discrete links between two devices. HDMI introduced sequences of links of up to four devices. Network technologies enable arbitrary communication among any number of devices. Networks are not traditional in A/V systems, though as Internet Protocol technologies make inroads, they are beginning to be a factor.

Streams

A stream is an ongoing flow of data that transports audio and/or video content. It flows from a source to a sink, across one more links, of the same or different technologies. It is an association between the source and sink devices, and the source context (e.g. channel, time in media, or playlist name).

Some A/V systems can have multiple streams operating at the same time. Examples are when one source is being recorded and a second source is displayed on a TV, and “picture-in-picture” features that show two streams on the same screen. Stream topologies are becoming more complicated with ability to replicate video content to a second TV, and split off audio from a video stream to a separate device.

A stream may contain typical TV or movie content, or a content navigation screen. Streams can also be static (e.g. when content is paused), audio-only, or audio with peripheral video content (e.g. artist and song title).

Sources

A source is the initial appearance of a stream within an A/V system. Examples include local static media (e.g. DVDs), recorded media (on DVRs), remote content (e.g. from broadcast or service providers), or locally created content (from security cameras, PCs, or game consoles).

Sinks

A sink is the final destination of a stream, with the most common sink a television display with integrated speakers. A “display” communicates media to a user. A TV’s main component is its visual display; speakers are “audio displays”.

Intermediate Device

Some devices are sometimes or always only a “pass-through” of content and neither a source nor a sink. An A/V receiver is the most common example of this. Some devices can be sources, intermediates, or sinks at different times or for different purposes.

Control Signals

Control signals are data not part of a media stream itself, but which actively determine or describe the content being displayed. Examples include infrared data from remote controls, and control paths on digital links such as HDMI. High-end A/V components sometimes include interfaces for control signals, some of which use RS-232 interfaces; these enable a central controller to direct the detailed operation of other devices.

Power Control in Current Systems

Today’s A/V systems are generally very good at accomplishing their primary function of displaying media content for people to enjoy. However, control of power states on these devices is frequently cumbersome and/or confusing. For a variety of reasons, devices are often left fully on when not needed, wasting energy. Users are annoyed and distracted by manually powering up devices when needed.

The core issue this paper addresses is the matching of content wanted (and actually consumed) by people, with devices being active as necessary to deliver that content. Energy efficiency is best served by doing this as precisely as possible, with the least burden imposed on users.

An Improved A/V Architecture

This section describes principles for a new A/V power control architecture, which builds on existing technologies, but makes key changes and additions. The core concept only addresses how devices should behave in an optimal future with all new devices; how we manage the transition is considered later. Power control is primarily about transitions – how a device knows when its state should change. This section covers principles fundamental to the new approach, issues in changing power state, what this means for how devices behave, and user interface concerns.

Creating this new architecture is part of an ongoing project; earlier phases [2] reviewed characteristics and features relevant to power control, first of current A/V devices [3], and then for communications technologies [4]. The most recent phase [5] explored how the concept of sleeping streams would affect the operation of devices and protocols, by assessing use cases for device and stream operation, resulting needed device behaviors, and particular issues that require special attention.

Technology architectures cannot be proved to be optimal nor can they be derived from data in the way that many physical principles can be. Rather, they can simply be described and explored, as well as compared to other technologies for their capabilities and burdens (such as complexity). An alternative architecture to the sleeping stream one would include a central control device that would manage the power state of all A/V devices in a local network. This approach appears to be more costly, less effective, and more likely to fail than the sleeping stream approach. One could require that all devices participate in a single network protocol. This is unlikely, at least for the foreseeable future, and in comparison to the sleeping stream model, seems unnecessary. Thus, this scheme is proposed as the best solution to the problem, though alternatives are welcomed to test this hypothesis.

Starting Principles

This project did not arise in isolation; rather, it builds on lessons learned over the last several decades in how to manage digital devices in a network context. A/V devices are beginning to adopt many conventions and technologies from computers in how they communicate and behave. Examples include the change from using only data links, to increasingly making use of network connections (that are digital and bidirectional), and the ability of some A/V devices to perform multiple tasks simultaneously. Thus, lessons from how PCs are power-managed are likely to apply to A/V devices. Key among these are:

- Use a three-state power state model (not two) and make it clear in the user interface.
- Maintain network connectivity in sleep.
- Ensure that power management is as automatic as possible.
- Keep delays on device wakeup short (a few seconds at most).

These lead to several conclusions about devices; they should:

- Be aware of the power state of other A/V devices in the local network.
- Be aware of the functional state of other A/V devices in the local network.
- Mostly toggle between on and sleep.
- Be quick to wake, and (relatively) slow to go to sleep.

In constructing the use cases for this project (see below), a design principle was that users care most that devices are on and available when they are wanted so that it is important to minimize or avoid forcing people to power up devices manually (e.g. with its remote control, its power switch, or with a different remote programmed to do so). Devices need to always wake if they need to (or even just might need to), to maximize user convenience and match expectations. If they fail to do this, users will likely disable power saving features leading to much energy waste.

Most people value convenience and a system of powering-down devices that is automatic, and an automatic system will engage considerably more often than manual-based ones will.

These principles, and consideration of how people do today, and will in the future, use A/V devices led to the sleeping stream concept and associated content.

Analysis

This project phase began with an assessment of abstracted use cases of device operation that affect power states in the context of sleeping streams. There is a minimum set of use cases that covers the needed complexity, and beyond this, additional cases only raise issues already covered by the basic set. For example, cases with more than three devices simply have several intermediate devices, each of which has the behaviors of the intermediate in the three-device case, so that no cases with more than three devices are needed. Table 1 shows an example use case with two devices and one content stream. The START and END lines show the power state of each device, while the intermediate lines show the sequence of actions taken by individual entities. If the TV had already been on at the start of the process, the resulting sequence would be effectively the same, since waking the stream would not require step 3 to wake the TV but would be otherwise identical. This is an example of an alternate use case that would add no new device behaviors.

Table 1 refers to a “fixed” stream. This is a content stream of finite duration, such as a movie. Other streams, such as broadcast media, are continuous and have no definite limit. Device behaviors will sometimes differ between these two stream types. The study [5] presented five one-device use cases, nine two-device cases, and two three-device cases. It also considered two cases in which the stream was addressed directly, rather than one of the devices, as well as two example cases of “failure” when a device does not behave as intended and so the resulting action is compromised. It is abnormal cases such as those that include failure that can create the most complexity for manufacturers and the most difficulty for users, so it is helpful when the core technology is as simple as possible.

Table 1. Use case example: DVD player powers up

Step	DVD	TV	Stream	Comments
START	Sleep	Sleep	Sleep	
				DVD power-up command (manual or internal timer) or manual play command
1	Wake			
2			Wake	DVD wakes up last stream it participated in
3		Wake		Stream involves TV so TV must power up
4		Input		Change Input (If necessary)
5	Play		On	Only after both devices fully wake (only applies to fixed streams)
END	On	On	On	

Note that this architecture does not specify the mechanisms by which the communication occurs that enables these sequences of actions to occur – just the result. Other standards — extensions to existing ones — provide the mechanisms.

In addition to device (and stream) power states and actions that entities engage in, the use cases also contain other elements. One is user action, which can be direct (e.g. use of a remote control), or indirect (e.g. change of state in an occupancy sensor). Another is a delay timer, such as a period of time of no activity that might initiate a power-down sequence, or the length of time that a warning is displayed of an imminent power-down. A third type of action is unexpected failures of a device or communications link.

The use cases outlined were evaluated to identify standard behaviors that devices should implement. Table 2 shows a portion of the result – behaviors that a source device must have. There are corresponding tables for sinks, intermediates, streams, and failure modes. While the use cases in the full report [5] on first glance suggest significant complexity, the behaviors show that in fact the system is based on a modest number of rules. The relative simplicity of the system is good for the design of communication protocols needed to accomplish it, the design of products which implement the protocols and behaviors, and for the ordinary human beings who must use the resulting products and systems.

Table 2. Source Behaviors

On

Notified stream to go to sleep
- Pause stream if fixed
- Go to sleep

Source paused
- Wait for X time
- Tell stream to sleep

Fixed Streams only

Fixed content ends
- Menu for X time
- Tell stream to sleep

"All devices ready" signal from stream
- Play content

Sleep

Powered up
- Wake self
- Wake stream

In the course of the analysis, it became clear that several temporary transitional states were needed for such a system to work. Figure 3 shows a proposal for how the three basic stream states and three transitional states might relate to each other; "GTS" is going-to-sleep. Devices commonly have transitional states internally but the system architecture may not need to represent them externally in the way that it needs to be cognizant of detailed stream states.

Content streams today are usually simple, involving no more than three devices all arranged in a linear fashion. However, it is quite possible, and even likely, that future streams will include branches into or out of the core stream as well as being longer than three devices. Figure 4 shows this graphically, but it is not apparent that such complexity needs to change the basic operation of streams sleeping and waking, except for how failure modes are treated.

Figure 3. State diagram for stable and transitional stream states.

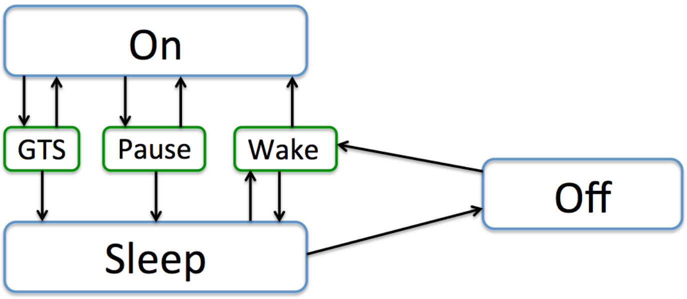

Figure 4. More complex stream topologies

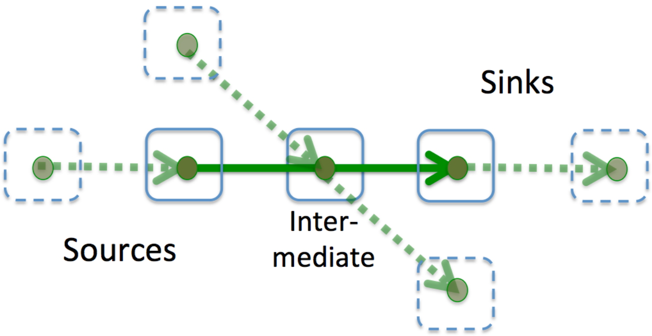

Sleeping Streams

Media content streams are the core of A/V system functionality and need to evolve to create new capabilities and to optimally support power control. Computers have three basic power states, and since applications are contained within a single device, they go to sleep with the device, and the operating system informs applications when sleep and wake events happen. In A/V devices, streams are the analog of applications, being the basic unit of activity. Since streams inherently involve multiple devices, power management of streams can be considered separately from the power state of the involved devices (that is, while stream state influences device state and vice-versa, they are often not the same). In general, the three-state model has advantages over its two-state counterpart.

A "sleeping" stream has its representation maintained within devices and networks, but no media content is communicated. This enables new functionality to more easily match the power and functional state of devices to what is desired by the user. A sleeping stream is more available than a stream which does not exist. Creation of a new stream requires a series of link and inter-device negotiation steps, requiring time and user effort. In contrast, waking a sleeping stream requires many fewer steps. A sleeping stream is distinct from one which is active but just paused, since a paused stream needs to be able to resume with no delay at all and should always contain a static image of the pause point. Pausing usually is for only seconds or minutes. Sleeping can extend for months or years. While a device can be turned off then turned back on again, the alternative to a stream being on or asleep is for it to be deleted. Once deleted, it must be recreated from scratch for it to exist again.

As with sleeping devices, people will need to be aware of the sleep state of a stream to understand its functionality, as it may be useful in how devices are used. For example, the video stream from a security camera could be set as a sleeping stream and so have authority/capability to wake up a specific television (or group of displays) when it becomes active. By contrast, a mobile phone entering a space that does not have a relationship with the TV might need to have the TV powered up first before it can negotiate access to the display.

Another relevant analogy for streams in a network context is the Transmission Control Protocol (TCP), widely used in Internet communications [6]. It enables reliable bidirectional data transfer between two end points. Before any data can be transferred, a negotiation takes place between the two end points to agree to open the data connections, and the device that gets the request from the other can refuse it. Either side can close the connection at any time. TCP connections do not have the concept of a sleep state (although they probably should), but if an extended period of time passes with no data being sent, then either device can send a "keep-alive" packet with no stream data to assure the other side that it is still there. Since TCP includes acknowledgment packets, presence of both devices on the network is confirmed any time data or keep-alives are exchanged. The energy burden of occasional packets is minor, so long as doing this does not require a higher device power state. Thus, the existence of a potential flow of data is decoupled from the actual flow of data.

As the use case and behavior analysis showed, the combination of sleeping devices and sleeping streams provides needed flexibility for a variety of current and emerging usage scenarios while keeping the complexity of the system manageable for both devices and the people that use them.

An attribute of the sleeping stream concept is that when a new stream is to be created, all of the involved devices must be fully awake. This makes the complexity that a sleeping device must implement much less than otherwise. For example, security and authentication present significant challenges and these only are dealt with by devices that are fully on.

User Interfaces

Past experience with power control of electronics has shown that unclear or inconsistent user interfaces are a barrier to saving energy. In addition, standardizing user interfaces has a very low to no manufacturing cost, and improves the overall user experience. One need is to clearly embody the 3-state power model into A/V devices, in hardware (including indicator lights), and software controls [7,8]. Another need is to ensure that user interfaces for device control include device power state information, so that when devices available are presented, there are subtle distinctions made between those that are fully on and those that are asleep (off devices will generally not be visible at all). This distinction can help alert users when there are functionality differences between sleeping and on devices. It can also help a user to diagnose problems, e.g. to see when devices are fully on that don't need to be (many users do not commonly look at most devices, having them in cabinets, closets, or elsewhere).

Inertia is a powerful force in user interface design, and power control is not a feature that is likely to be a driver of purchase decisions, so manufacturers have little incentive to focus on it. However, it is clear that the energy savings benefits from improved user interfaces is orders of magnitude larger than the costs required to deploy them.

A further user interface challenge is how to represent content streams other than those actively being viewed. It remains to be seen what overall principles and approaches manufacturers will bring to this so it is premature to comment on it from the energy perspective. However, a clear visual distinction between streams that are on or asleep is needed, and this can be standardized once it becomes clear what that is (a possible approach is to dim or 'grey-out' a textual or iconic representation of a sleeping stream).

Other Considerations

In the analysis, we considered a variety of issues that the architecture, protocols, and devices must cover. The following are the major issues considered; full detail in the main project report [5].

Named streams. For device-device communication as well as device-human communication, streams will need to be distinguished from each other, which requires some mechanism for identification. The simplest way to do this is with a human-readable name. Such names will need to be unique within a local network.

Occupancy. As the purpose of audio/video streams is to communicate information to people, the occupancy of a person or people in a space can be critical information to know when to begin or end displaying a stream. Mechanisms to acquire and distribute occupancy information among devices will be a key technology innovation. This can also be extended to other uses of energy such as lighting and climate control.

Failures. When a device or communications link fails, the interaction between people and devices can be the most intense and important. Thus, how devices respond to failure conditions is critical to their user amenity and ultimately energy savings.

Multiple streams in a network. Local networks may contain multiple simultaneously active content streams, with some devices involved in more than one.

Multiple streams per device. While today collections of A/V devices usually only involve a single content stream at one time, in future many of these will be capable of, and commonly implement, more than one simultaneously. This will require adaptations to technologies, products, and user experience.

Multiple sinks and/or sources. Streams today usually involve a single source device and a single sink (an exception is when an A/V receiver sends audio to speakers independently of the video signal sent to a TV). As the availability of sources expands, we can expect more ability to merge multiple sources into a single stream. In addition, as the cost of displays drops, and the ability to direct data to multiple displays increases, it will be attractive for people in homes (and other use contexts) to replicate a content stream across multiple displays. Both of these will require the same sorts of changes that multiple simultaneous streams do.

Creating and maintaining streams. With persistent existence, there will necessarily be some overhead for both people and devices in managing streams. The burden of this needs to be kept as low as possible while providing needed capability and flexibility.

Emergency broadcasts. As content streams become networked, the possibility of using them for disaster notifications (natural and otherwise) will become apparent and implemented. This will require specific network availability of devices, authentication needs, and user preferences for this purpose.

The Transition

Even when a well-functioning system is implemented in all new devices, we will still have a huge stock of existing “legacy” devices that the new devices will need to interoperate with. Systems with legacy devices will use more energy and lack usability advantages; effort will be needed to assure that problems in both areas are minimized. New devices will need to detect when they are interacting with legacy devices and so adjust their behaviors. Some existing technologies have command/control mechanisms useful here. A new device could use these to manage the power state of legacy devices it is connected to. New devices could also provide user direction to manually manage legacy device power state when no other mechanism is available.

Next Steps

This architecture described here is being presented to the consumer electronics industry for consideration. If accepted, existing protocol standards will require extensions, in particular, adopting the three-state power model for devices and streams, and always exposing power and functional state to other devices on the local network. If accepted, the following steps are anticipated to result in products that implement it.

- Create a standard that describes the overall scheme of sleeping streams and associated details that could be forwarded to standards organizations for consideration. This should enable product designers to understand what to do and why, how to present this to users, and how to adapt the system to new circumstances.

- Review communication technologies to determine gaps between what they do today and what the architecture describes.
- Fill those gaps by amending technology standards.
- Put the technology into new products. Some features can be put into products even before the standards development is finished.
- Create a detailed summary of recommendations for how to deal with legacy products.

Critical to the process will be to get critical review from manufacturers of A/V devices, and the standards committees for relevant interface technologies. If a critical mass of these individuals and organizations do not support the system, it will not succeed. Manufacturers need to be assured that user amenity is the top priority.

For communications standards, the principal one needed to accomplish this in use today is HDMI [9]. Apart from providing the needed data path to send video data, HDMI also has features for content protection, addressing piracy concerns for many companies. HDMI is most commonly understood as a point-to-point mechanism, but it does facilitate a tree of devices with a single display at its root. It is possible to directly embed the notion of a sleeping stream within the core HDMI protocol, or implement it at a "higher layer" over the Ethernet channel present in newer versions of the standard. The UPnP protocol may also be critical.

It is also likely that one or more standards will get wide use for transporting A/V streams over Ethernet and Wi-Fi. These standards will also need to implement sleeping streams.

Considerations for Energy Policy

Energy policy can advance this technology in time and could make a difference in it being created at all. Past examples have shown that the Energy Star program has done this for communications technologies, and the nature of information technology suggests that this will be more possible and needed in future.

Many energy policy programs have a goal to be "technology neutral" to not favor or require technologies that may unfairly advantage specific companies. This is sometimes extended to not wanting to require network interoperability technologies. This is a mistaken application of the technology neutrality principle, as open standards do not favor specific companies but are rather available to any company. In addition, some energy savings opportunities are only available if companies implement specific technologies so that failure to embrace them forecloses some potential savings.

Conclusions

A/V device control is a complex topic with diverse devices, interfaces, technologies, system construction, and usage models. As power control has not been a high priority for the industries involved, it is not surprising that it is not well articulated in current technologies and devices systems. Digital technologies not only bring functionalities not available previously, but also offer the possibility of a more robust system for power control. A new power control architecture seems possible to construct and implement in technologies and products. The bigger challenge will be to improve products in the interim that have some legacy analog interfaces which inherently complicate power control.

Acknowledgement

This work was supported by the California Energy Commission's Public Interest Energy Research program and by the Assistant Secretary for Energy Efficiency and Renewable Energy, Building Technologies Program, of the U.S. Department of Energy under Contract No. DEAC02-05CH11231.

Many thanks to Gari Kloss for her invaluable contributions.

References

[1] Urban, Bryan, Verena Tiefenbeck, and Kurt Roth. 2011. Energy Consumption of Consumer Electronics in U.S. Homes in 2010. Final Report by Fraunhofer Center for Sustainable Energy Systems for the Consumer Electronics Association.

[2] Nordman, Bruce, Audio/Video Inter-Device Power Control, ACEEE Summer Study on Energy Efficiency in Buildings, 2012.

[3] Nordman, Bruce, 2012. Existing Product Assessment and Conclusions Memo, interim deliverable for this project, June 15, 2012. http://nordman.lbl.gov/avcontrol

[4] Nordman, Bruce, 2012. Summary of Analysis of Communication Link Technologies, interim deliverable for this project, September 21, 2012. http://nordman.lbl.gov/avcontrol

[5] Nordman, Bruce, and H.Y. Iris Cheung, Use Cases, Analysis Method, and Candidate Device Behavior Model Report, December 21, 2012. http://nordman.lbl.gov/avcontrol

[6] Postel, Jon. 1981. Transmission Control Protocol, Internet Engineering Task Force, RFC 793.

[7] Institute for Electrical and Electronic Engineers. 2004. Microprocessor Standards Committee of the IEEE Computer Society. P1621, Standard for User Interface Elements in Power Control of Electronic Devices Employed in Office/Consumer Environments. (approved December, 2004; reaffirmed December, 2009).

[8] Nordman, Bruce. 2003. The Power Control User Interface Standard — Final Report. prepared for the California Energy Commission, Public Interest Energy Research Program, LBNL-52526.

[9] High Definition Multimedia Interface LLC. 2010. HDMI Networked Products and Consumer Electronics Control (CEC), HDMI Version 1.4a, March 2010. http://www.hdmi.org/manufacturer/specification.aspx